HAMILTONIAN COMPLEXITY

TOBIAS J. OSBORNE

Abstract. In recent years we've seen the birth of a new field known as *hamiltonian complexity* lying at the crossroads between computer science and theoretical physics. Hamiltonian complexity is directly concerned with the question: *how hard is it to simulate a physical system?* Here I review the foundational results, guiding problems, and future directions of this emergent field.

1. Introduction

In the natural sciences complex behaviour is usually exhibited by large numbers of interacting particles moving according to simple underlying physical laws. *Many body theory* is the area of physics aimed at explaining these systems, a task often requiring the study of equations comprising large numbers of variables. Problems involving many variables like these are a recurring theme throughout the natural sciences and there are several approaches to understanding their properties. The program that concerns us here is that of *computational complexity theory*, which is the branch of computer science and mathematics focussed on the task of classifying their *computational difficulty*.

There are striking parallels between the central objects of study in many body theory and computational complexity theory. For example, the *hamiltonians* of many body physics bear some superficial resemblance to the *propositional formulae* studied in complexity theory. It was quickly realised that this analogy runs deeper: the equilibrium states of many physical systems directly correspond to satisfying assignments of variables in certain families of propositional formulae. Following this it made sense to apply the apparatus of computational complexity theory, aimed at answering the question "how hard is it to decide a propositional formula?", to the study of many body physical systems. Although these analogies were explored in some detail during the 80s and 90s it was only with the advent of quantum information theory that we have witnessed an intensified cross-fertilisation between computational complexity theory and many body physics, leading to the birth of *hamiltonian complexity*.

The central objective of hamiltonian complexity is thus to answer the question: "*how hard is it to simulate a physical system*?" Computational complexity theory provides us with the apparatus to quantify what is meant by "how hard" and many body theory, in particular, condensed matter physics, provides us with a rich family of strongly correlated systems to study.

Hamiltonian complexity is interesting precisely because there are families of systems that could be, in principle, engineered and whose physical properties are hard to simulate. The canonical example here is due to Barahona [15] who described an Ising model in a nonuniform magnetic field whose *ground-state energy* is hard to calculate. By "hard" we mean that the computational resources required to estimate it grow too quickly with the system size, so that no computer can carry out the task in a reasonable time. Further, the origin of this hardness isn't just a lack of imagination: no matter how clever you are there is no way to efficiently extract predictions about these systems[1]. This peculiar situation is initially disturbing: if one cannot solve the model in a reasonable time, even with access to a supercomputer, then how can the model be said to have *predictive* content? Could it be that there exist systems whose behaviour is simply too hard too predict? This is reminiscent of chaos theory, but the origin of the lack of predictability is rather different.

So far this discussion applies only to classical systems. However, it was realised that *quantum computers* would be capable of simulation tasks exceeding their classical counterparts. Indeed, a quantum computer can

Date: June 30, 2011.

[1]This is because merely the *existence* of an efficient method to compute the ground-state energy of Barahona's model would serve as a counterexample to the $\mathsf{P} \neq \mathsf{NP}$ conjecture. I'll say more about this later. Obviously, for this review, I am adopting the position that this conjecture is true. If it were false then this would have some spectacular consequences for the natural sciences...

simulate the real-time dynamics of naturally occurring many body systems, both bosonic and fermionic [76]. With this observation came the hope that quantum computers would also be able to simulate the *equilibrium* dynamics of, e.g., Barahona's model, and much effort was expended on the task of developing such a quantum simulation algorithm. Unfortunately – or fortunately, depending on your point of view – it became increasingly evident that such quantum algorithms couldn't always succeed [1]. Concomitant with these investigations has been a parallel effort at classifying the *quantum complexity* of computational problems, culminating in the construction of quantum models which at first sight appear to be simulatable with a quantum computer, yet whose properties are, in a specific sense I'll describe later, "even more" intractable than their classical cousins.

The puzzling observation that there are systems whose equilibrium properties are hard to simulate has a relatively simple resolution: it turns out that the assumption that the system ends up in an equilibrium state modelled by the Gibbs distribution is simply *wrong*. This isn't really surprising, and is familiar in the context of spin glasses [16]: when a system has access to a large number of metastable configurations it is overwhelmingly to end up in a metastable state rather than the unique lowest-energy state. Indeed, Barahona's model, and its cousins, exhibit precisely this property. Thus, if you were to engineer a "difficult" system such as Barahona's, it would take an extremely long time for it to equilibrate to its ground-state, and any classical or quantum simulation algorithm for its ground-state properties would get stuck. Thus we've learnt that the manifold of states accessible by the natural dynamics of a physical system is a tiny subset of the full configuration space of the system. By understanding this corner of configuration space better we can design better variational classes of wavefunctions to best parametrise naturally occurring states.

The outline of this review is as follows. We begin in §2 with a brief discussion of computational complexity. In §3 we discuss the central question of hamiltonian complexity, *the simulation problem*, and review its origins and motivations. In §4 we review a large body of literature focussed on the presentation of physical systems that are *hard* to simulate. Finally, in §5, we review the literature focussed on the converse situation of understanding what physical systems *can* be efficiently simulated.

2. Complexity theory

In computational complexity (see, e.g., [81] or [10] for a textbook treatment) one categorises the difficulty of computational problems by organising them into *complexity classes*. To understand which class a problem belongs to you need to study how many basic computational steps (e.g., addition, subtraction) are required to solve the problem for larger and larger inputs. It is typical that the larger the input to the problem is the more computational steps are required to solve it. For example, consider the problem of sorting a list of n numbers in descending order. As the input size n gets larger it takes, *in the worst case*, no more than $O(n\log(n))$ steps to do this. Another example: consider the problem of multiplying two $n \times n$ matrices. This takes, in the worst case, no more than $O(n^{2+\epsilon})$, $\epsilon < 1$, operations to complete [29]. In computational complexity one doesn't care so much about the exponent on the n, rather, that there *is* some finite exponent $\alpha < \infty$ so that, in the worst case, no more than $O(n^\alpha)$ operations are required to solve the problem.

The class of all *decision problems* – i.e., problems with a "yes" or "no" answer – solvable using a number of basic computational steps scaling no worse than a *polynomial* in the input size n is called P. The exact polynomial rate needed to solve a problem in P may differ from computer architecture to architecture (the actual definition of P doesn't depend exactly on the definition of *computer*), but the fact that the rate is polynomial doesn't. This is an admittedly coarse classification because for all practical purposes a problem requiring n^{1000} steps to solve is intractable for large instances. Interestingly, for many problems in P the exponent seems to be low.

Another fundamental complexity class is NP which is, roughly speaking, the class of decision problems whose solution may be *checked* using a polynomial number of basic computational steps. At first sight this innocuous definition doesn't hint at any hidden intractability. However, it turns out that NP contains extremely complex problems. Firstly, it is worth noting that $\mathsf{P} \subset \mathsf{NP}$ since the solution to any problem in P can be checked using polynomial steps (just solve it yourself). To convince yourself that NP might contain more problems than P consider the following situation. Imagine you have some function f of n bit values $x_1, x_2, \ldots, x_n$, $x_j \in \{0, 1\}$, and you want to find its minimum value. (One can easily convert a problem such as minimum finding into a decision problem by solving the related problem: is the minimum of f less that c? If you can solve this problem then you can use a binary search to find the minimum value up to some

precision.) It is convenient to imagine that the bit values $\mathbf{x} \equiv (x_1, x_2, \ldots, x_n)$ represent the orientations of a collection of n spins and that f represents an *energy function* of $\mathbf{x}$. Now, as long as it is possible to compute $f(\mathbf{x})$ using only a polynomial number of operations, one can see that this problem of finding the minimum is *in* NP. (That is, if someone gave you the configuration which is claimed to achieve a value of f less than some c then it is easy to check that this is the case.) But is it in P?

There is one obvious strategy to find a value of f less than some desired c: simply search over the space of all bit strings $\mathbf{x}$. However, this task is prohibitively expensive (it takes $O(2^n)$ steps in the input size n). Perhaps there is a better strategy? Unfortunately not, unless f has some special structure. For example, if f is comprised of a sequence of random interactions – each coupling collections of many spins – the system will be *frustrated* and f will possess many local minima, ruining any local search heuristic.

The canonical example of such a frustrated f is provided by the 3SAT problem: here f is given by a Boolean expression written as a conjunction (i.e., the "OR", denoted $\vee$) of a collection of *clauses* C_j, each of which is the disjunction (i.e., the "AND", denoted $\wedge$) of three *literals* (i.e., a Boolean variable x_j or its negation $\overline{x}_j$). An example of an allowed clause is, e.g, $C_j(\mathbf{x}) = \overline{x}_{j_1} \vee x_{j_2} \vee x_{j_3}$. There are 8 possible choices for which variables can be subjected to flips and all are considered to be possible. You should think of which variables are actually subjected to flips in each clause, and which variables appear in each clause, as being chosen by an adversary. The function f corresponds to a (classical) hamiltonian for a system of n spins with general 3-spin interactions via

$$f = \bigwedge_{j=1}^{m} C_j(\mathbf{x}) \longleftrightarrow H(\mathbf{x}) = -\sum_{j=1}^{m} h_j(\mathbf{x}), \tag{2.1}$$

where each of literals x_j are interpreted as numbers via TRUE = 1 and FALSE = 0 and the interaction terms h_j are found by replacing $\vee$ with the standard product operation, e.g., $h_j(\mathbf{x}) = \overline{x}_{j_1} x_{j_2} x_{j_3} \longleftrightarrow C_j(\mathbf{x}) = \overline{x}_{j_1} \vee x_{j_2} \vee x_{j_3}$. A configuration of zeros and ones achieving the minimum value of H corresponds to a *satisfying* assignment of f where all of the clauses evaluate to TRUE. If zero is regarded as a spin up and one is regarded as a spin down then H can be thought of as the hamiltonian for a (possibly frustrated) spin system.

It is not impossible that there *is* a general-purpose efficient algorithm to find the minimum value of H, but this would have fabulous consequences because of a fundamental result in computational complexity theory known as the *Cook-Levin theorem*: the 3SAT problem is *complete* for the class NP. The theorem says that if you had an algorithm which could solve 3SAT using a polynomial number of steps, *in the worst case*, then you could solve *any* other problem in NP efficiently. Any such problem in NP with the property that an efficient algorithm for it would yield an efficient algorithm for all of NP is known as an NP-*complete problem*. A problem is known as NP-*hard* if it as at least as hard as the hardest problems in NP. There is no requirement that such problems be in NP; indeed, they may not even be decision problems. The class NP is now known to contain a vast number of problems, many of which have considerable practical relevance. However, since no efficient solution algorithm has ever been found, nor are there serious proposals, the possibility of P = NP is considered to be extraordinarily unlikely.

Interestingly, there are some surprising examples of systems that appear, at first sight, to be intractable, yet which admit efficient algorithms for their physical properties. A good case in point is that of a *ferromagnetic* Ising model – where the spin-spin interactions are always positive – with arbitrary interaction graph. Thanks to Barahona's example (which pertains to an Ising model with both antiferromagnetic and ferromagnetic interactions), one might be tempted to assume that the equilibrium properties of this model cannot be simulated efficiently. However, there is a clever polynomial-time algorithm which samples configurations according to its Gibbs distribution [51], to a required precision.

With the advent of quantum compution the landscape of computational complexity theory has been fundamentally extended: it is natural to introduce a new class of problems extending the classical class[2] P, which are easy for a quantum computer to decide. This natural class is known as BQP and contains several problems thought to lie outside of P, e.g., factoring. The classification of the computational complexity of problems in the presence of quantum computers is known as *quantum complexity theory*.

There is also a natural class of problems known as QMA, analogous to NP. To understand this generalisation it is convenient to view NP as a game played between two players: (M) a malevolent Merlin, with infinite

[2]Strictly speaking, the class BQP extends the class BPP, which are the decision problems that can be decided with a bounded probability using a polynomial number of arithmetic operations and coin flips.

computational power; and (A) Arthur, limited to using a standard classical computer with polynomial computational power. In this game Arthur asks Merlin questions, such as "does f have a satisfying assignment?" and Merlin accompanies the answer with a *proof* $\mathbf{x}$, which Arthur can use to verify the answer. In the case where f doesn't have a satisfying assignment there is no way Merlin can ever fool Arthur that it does. The class NP is then the class of all decision problems which can be decided via this game.

The quantum generalisation of NP is defined along the same lines, except that Arthur is allowed access to a *quantum computer* with polynomial computational power and Merlin is allowed to send quantum states as a proof. Arthur asks the same questions as before, but now Merlin is allowed to send a quantum state to convince Arthur that a given problem has a "yes" answer: Arthur *accepts* that a given problem has a "yes" answer if, after running a quantum verification algorithm on Merlin's quantum state, he measures a 1 and *rejects* if he measures a 0. We say a decision problem is in QMA if Arthur always accepts a "yes" instance with probability greater than 2/3 and always rejects a "no" instance with probability less than 1/3.

A remarkable result of Kitaev (discussed below) establishes a Cook-Levin-like theorem for QMA, and provides a complete problem for it. The main intuition to take away here is that the class QMA contains "harder" problems than NP.

3. The simulation problem

Typically the way physics has done is as follows: (1) formulate a theoretical model; (2) (approximately) solve the model to make a prediction; then (3) experimentally test the prediction. If the experiment confirms the prediction then declare success. If not, then go back to the drawing board. Each of these three steps typically requires considerable ingenuity, insight, and resourcefulness to carry out. However, of these three steps, the second has a somewhat algorithmic flavour: if one only threw enough computational resources at the model then surely it could be solved? This echoes Hilbert's *Entscheidungsproblem*: perhaps for any physical model, when sufficiently clearly formulated, there exists a procedural algorithm to extract predictions, at least up to some acceptable error margin? The analogy with the Entscheidungsproblem here has been specifically chosen to foreshadow the answer: *no*, there is no such general efficient method.

Hamiltonian complexity has been developed to answer precisely the question: *how hard is it to simulate a physical system?* To attack this vaguely phrased question it is convenient to introduce a decision problem to capture its essential elements. This problem, called the *simulation problem*, is described as follows. Firstly, we need to choose a physical system. In all cases we assume the system (usually quantum) is modelled with a hamiltonian H. A crucial ingredient in the simulation problem is that we don't just have one system but rather the specification of a *family* $\{H_n\}_{n=1}^{\infty}$ of physical systems (usually indexed by the number, n, of particles involved). This is fundamental because in order to answer the question "how hard is it to simulate H_n?" using the tools of computational complexity we are going to quantify the hardness using the *scaling* $f(n)$ with n of the number of arithmetic operations required to simulate the system.

Next, "simulate" is far too vague: we also need to specify what state the system is in, i.e., thermal equilibrium or dynamically evolving? Additionally we need to say *what* we wish to simulate because for systems comprising n particles there are typically exponentially many possible quantities which are, in principle, observable, but most of these are experimentally inaccessible. Thus the final ingredient of the simulation problem is the specification of a family $\{A_n\}_{n=1}^{\infty}$ of observables to simulate. Finally, it is unlikely that, constrained as we are by the finiteness of machine precision, we can make perfect predictions for A_n. Therefore we specify two expectation values α_1 and α_2 which our predictions *must* decide between. With these qualifications in mind we now write out the simulation problem.

Problem 3.1 (The simulation problem)**.**
Input:

- A specification of a family $\{H_n\}_{n=1}^{\infty}$ of hamiltonians, a specification of an initial state $\rho_{0;n}$, a family of observables $\{A_n\}_{n=1}^{\infty}$, a (possibly complex) time t, two proposed expectation values $\alpha_1 < \alpha_2$, and an integer n.

Output:

- "yes" if

$$\langle A_n \rangle \leq \alpha_1, \tag{3.1}$$

or

- "no" if

$$\langle A_n \rangle \geq \alpha_2, \tag{3.2}$$

where

$$\langle A_n \rangle = \mathrm{tr}(\rho_n A_n), \tag{3.3}$$

and

$$\rho_n = \frac{(e^{itH_n})^\dagger \rho_{0;n} e^{itH_n}}{\mathrm{tr}((e^{itH_n})^\dagger \rho_{0;n} e^{itH_n})}. \tag{3.4}$$

Note that, because we allow t to be possibly imaginary, e.g., $t = i\beta$, the trace factor in the denominator may be nontrivial. From now on we drop the subscript n on H and ρ and leave the dependence implicit.

Two special cases of the simulation problem have received considerable attention in the literature, namely, the *real-time* case where t is real, and the *ground-state* case where $\rho_0 = \mathbb{I}/\,\mathrm{tr}(\mathbb{I})$ and $\rho = \lim_{\beta\to\infty} e^{-2\beta H}/\,\mathrm{tr}(e^{-2\beta H})$ which is discussed in the next section.

As we'll see, much effort has been expended on understanding the simulation problem in the case where the observable A is the hamiltonian H of the system itself. I personally think that there is much to be gained by studying more general observables A: not much known about the complexity of the simulation problem for situations of immediate physical relevance, e.g., when A is *bulk* or *extensive* quantity such as the magnetisation of a quantum magnet. But surely this should be an easy problem to simulate? After all, thermodynamics tells us how bulk properties should behave. Obviously it all depends on what tolerances $\epsilon = |\alpha_1 - \alpha_2|$ you choose: there should be a smooth transition in the complexity of the simulation problem from hard to easy as ϵ is varied between, say, $\epsilon = O(2^{-n})$ to $\epsilon = O(n)$. Indeed there is some evidence of such a transition: for ϵ small enough the problem of estimating the magnetisation is essentially equivalent in hardness to the ground-state problem [37] and for $\epsilon \sim n$, in the case of certain *regular* lattice systems, there is an efficient approximation scheme to estimate the ground eigenvalue [14].

Another direction which has received considerably less attention is that of t complex, but not pure imaginary, nor pure real. This problem is of considerable physical importance as it corresponds to the situation where a system is dynamically evolving at some *finite temperature*, which is the generic situation in experiments where it is usually impossible to cool a system to its ground state. Looking further afield, it is also interesting to propose the study of the computational complexity of systems in the presence of *decoherence*. In both the finite-temperature and the decoherent cases there ought to be a transition in the hardness of the simulation problem from hard at low temperature/decoherence rate to easy for high temperature/decoherence rate.

4. What is hard

In recent years we've seen a string of counterintuitive results presenting realistic physical systems whose observable properties are *hard* to estimate. As previously mentioned, a foundational result was presented in a paper of Barahona [15] who proved that for the hamiltonian H of the *classical* Ising model in a nonuniform magnetic field the *ground-state* (i.e., $t = i\lim_{\beta\to\infty}\beta$) simulation problem is NP-hard for the observable $A = H$ with $\alpha_1 = 0$ and $\alpha_2 = 2$. Barahona's construction – pertaining to a classical model which, in our quantum context, simply means a model whose hamiltonian is diagonal in a product basis – already presents a family of realistic systems whose physical properties are hard to predict: any efficient general purpose algorithm to simulate the ground-state energy of Barahona's Ising model would already be capable of deciding $\mathsf{P} = \mathsf{NP}$, and hence doesn't exist. It is worth commenting here that these results pertain to the *worst case* where the model is engineered adversarily. It is fairly likely that the *average case* complexity, where the interactions are not chosen adversarily, of a physical model will, in general, be quite different from the worst case.

The most important foundational result for hamiltonian complexity in the quantum setting is that of Kitaev [57] (for a review see [7]) who showed that the ground-state simulation problem where $A = H$ is a 5-local hamiltonian (i.e., a hamiltonian for a quantum spin system whose interactions involve no more than 5 spins at a time) with $\alpha_1 = 0$ and $\alpha_2 = O(n^{-3})$ is *complete* for QMA. Kitaev's proof exploited a clever trick inspired by Feynman generalising an idea exploited in the proof of the Cook-Levin theorem. This construction has been influential in all subsequent works and can be regarded as a truly quantum generalisation of the Cook-Levin

theorem. Thus QMA has the same property as NP: any efficient classical algorithm for the ground-state simulation problem for Kitaev's H would yield an efficient classical algorithm for any other problem in QMA. Since QMA contains BQP, this would imply that all quantum computations are classically simulable.

It may be reasonably argued that a 5-local hamiltonian isn't particularly natural, and this prompted researchers to explore exactly what kinds systems might have ground-state properties which are hard to approximate. Building on Kitaev's construction a sequence of results refining the original have emerged [55, 54, 75, 53, 77, 6] culminating in the surprising and counterintuitive result [6] that there exist even *one-dimensional* quantum spin systems whose ground-state energy is hard to simulate within a precision $\epsilon = O(n^{-\alpha})$, $\alpha \geq 3$. This remains true even when the system is translation invariant [36]. It is worth pausing a moment to appreciate the full implications of this result. In the classical case the one-dimensional ground-state simulation problem is in P because one can simply evaluate the partition function via a straightforward product of transfer operators. The results [6] of Gottesman *et. al.* conclusively establish that no quantum analogue of the standard transfer operator method exists in the quantum setting.

In the proof described in [36] it seems necessary that the local spin dimension needs to be rather large (at least more than, say, 8 or 9). It is an interesting question to study whether reducing the local spin dimension makes the problem easier. Some evidence this might be the case arises in the case where constraints are chosen randomly (see the discussion below concerning quantum satisfiability). Here it is known [67] that there is a transition in the *entanglement* structure of the ground state of one-dimensional quantum systems as the local spin dimension changes from 2 to 4. The case of 3-dimensional local spins seems special. Based on this I would be willing to conjecture that the ground-state simulation problem for one-dimensional chains of spin-1/2 degrees of freedom is in P. (Some of my colleagues disagree with me that this could be true. Any evidence either way would be quite interesting.)

While initial research focussed on the setting of distinguishable quantum spins, much remains true in the physically more realistic setting of interacting particle systems comprised of bosons or fermions. Along these lines the simulation problem for a family of systems known as *stoquastic*, whose hamiltonian's off-diagonal entries are negative, has been investigated [21, 63, 52, 24]. Such systems do not exhibit a "sign problem" because the coefficients of the ground state are nonnegative. (All bosonic systems in second quantisation with the standard kinetic+potential energy structure exhibit this property.) Similar results hold here: the ground-state simulation problem for stoquastic hamiltonians has been shown to be complete for the class *stoquastic*-MA which is, roughly speaking, a class generalising NP with a complexity intermediate between two probabilistic classes MA and AM extending NP. Additionally, the original QMA results developed above have been used to show that ground-state simulation problem is QMA hard for both fermions [91] and bosons [98].

A very closely related problem (indeed, dual to) to the simulation problem known as the *density operator consistency problem* has also received attention. Here the task is to decide if, given a list of reduced density operators, whether there exists a global pure state consistent with these reductions. This problem is now known [62] to be hard for QMA. The fermionic equivalent, the *N-representability problem*, of central importance in quantum chemistry, has also been shown to be QMA hard [64].

Parallel to this there have been several investigations exploring physically motivated specialisations, generalisations, and alternative settings of the original local-hamiltonian results. Notable results along these lines include studies of the complexity of state preparation via cooling [50] of local hamiltonian systems, and the difficulty of the task of simulating the density of states [92, 26] of a local hamiltonian.

There is one special case of the simulation problem for local hamiltonians which is of particular interest, namely, *quantum k-satisfiability*: here the task is to decide whether the hamiltonian is *frustration-free* or not, i.e., whether the ground state is *exactly* the ground state of each interaction term, or is far from such a state. This situation corresponds to the simulation problem where $A = H \geq 0$ with $\alpha_1 = 0$ and $\alpha_2 = 1/\text{poly}(n)$, i.e., the task is now to simply decide if a given state is *exactly* a ground state or far from one when given the additional promise that the hamiltonian is comprised of a sum of projection operators. This problem is arguably a more natural quantum generalisation of 3SAT. The same proof applies, as above, for local hamiltonians to show that this problem is hard in exactly the same settings, i.e., for one-dimensional translation-invariant quantum spin systems. The case of frustration-free *stoquastic* hamiltonians [24] is also rather interesting: here one can simulate the adiabatic evolution of such models classically.

The quantum k-SAT problem has turned out to be a fruitful basis for quantum generalisations of the *random satisfiability* problem, where constraints are chosen randomly. It is now understood that so-called random quantum k-SAT exhibits a transition [59, 60, 58, 67, 23], analogous to its classical counterpart [56], from a phase where there are almost always many possible frustration-free ground states to a phase where there are almost no frustration-free ground states.

A setting which has been the subject of considerable conjecture and interest is that of quantum spin systems with *commuting* interactions [25, 5, 89]. One might be forgiven for initially guessing that such systems are completely classical because all interactions can be simultaneously diagonalised. However, a little thought shows that this simultaneous diagonalisation isn't particularly helpful because the basis in which all the interactions are trivial may be highly entangled. Thus, the ground-state simulation problem may well be harder than its classical counterpart. Indeed, some evidence this might be the case has been recently obtained [5]. Deciding whether the simulation problem for commuting interactions is harder than NP, or in NP, is an intriguing open problem because it is clear that it requires new ideas going beyond the existing techniques.

4.1. **Guiding problem: the quantum PCP conjecture.** To date, in all of the constructions of quantum systems whose ground-state energy is hard to simulate, the precision $\epsilon = |\alpha_1 - \alpha_2|$, or *promise gap*, for which the problem is known to be hard scales as an inverse *polynomial* in the system's size. That is, the problem is known to be hard only when one requires that the energy is very precisely approximated. It is thus natural to ask to what extent the intractability of these systems is tied to this constraint, i.e., perhaps it is easy to decide if the ground-state energy is small or larger than, say, some *constant fraction* of the number m of interaction terms?

In the classical setting this question has a dramatic and counterintuitive answer: it is also complete for NP. This result, captured by the celebrated *PCP* theorem[3] [12, 11, 31], shows that even the task of approximating the ground-state energy to within a constant factor is hard.

The PCP theorem already establishes that there are physical systems, indeed classical systems, for whom the simulation problem is hard, even when $\epsilon = |\alpha_1 - \alpha_2| = cm$, with c some constant. However, it is entirely natural to suppose that that there should be a corresponding quantum equivalent of the PCP theorem which would say that the ground-state simulation problem is QMA-complete for $\epsilon = |\alpha_1 - \alpha_2| = cm$, with c some constant [3, 5].

Some progress towards a possible proof of the conjectured *quantum PCP theorem* has been made whereby an important proof element known as *gap amplification* – exploited in a recent proof [31] – has been generalised to the quantum setting, where it is now known as the *detectability lemma* [3]. However, so far, a fully quantum PCP theorem has resisted all attempts at a proof.

The quantum PCP conjecture is a fascinating guiding problem because it would require the development of many theoretical tools of independent interest. For example, it is considered likely by practitioners that a proof of the quantum PCP conjecture would aid in the design of error-tolerant architectures for *adiabatic* quantum computation. Additionally, proof techniques developed in the pursuit of the quantum PCP conjecture also have direct applications in the study of the local hamiltonians appearing in condensed matter physics. For example the detectability lemma has already found applications in the study of entropy/area laws for quantum spin systems (see below) [4]. A quantum PCP theorem would mean that there exist certain realistic hamiltonians for which no variational representation of the ground state with a concise classical description could reproduce the ground-state energy better than a constant precision.

At this point it seems that a proof of the quantum PCP conjecture would have an inductive structure whereby a very bad quantum PCP is iteratively improved. However, there is a stumbling block in any proof following this strategy, namely, the design of at least one, possibly very inefficient, quantum PCP. This is summarised by the following problem: show that a QMA-complete simulation problem for n spins can be encoded as the simulation problem for a local hamiltonian involving an *exponential* (or even *doubly exponential*) number N of spins, i.e., $N = 2^n$, or $N = 2^{2^n}$, with $\epsilon = |\alpha_1 - \alpha_2| = cN$, c a constant. In other words, if we don't care how many spins we use is it possible to improve the promise gap $\epsilon = |\alpha_1 - \alpha_2|$ *at all* while retaining the hardness of the problem? Historically this was one of the steps toward the classical PCP theorem. What I have in mind here is something like a quantum generalisation of the results of Håstad [40] on inapproximability of certain problems involving, e.g., solutions to linear equations.

[3]PCP stands for probabilistically checkable proof.

5. What is easy

In the previous section we saw that considerable effort has been expended on the task of constructing plausible physical models whose *ground-state* properties are hard to approximate. In this section I'll take a look at the flip side of the coin, i.e., understanding what physical properties are easy to approximate. To answer this question concretely requires the description of an *algorithm* which is *correct*, *converges*, and has a polynomial *complexity*. The design of such algorithms is truly an art form, but one informed by physical results.

The simple and striking conclusion we've drawn so far is that the natural dynamics of a locally interacting system doesn't always take the system an equilibrium state modelled by the Gibbs distribution. The question then arises: what states can a locally interacting system enter under natural dynamics? The key insight to draw on here is that the dynamics of a locally interacting collection of particles cannot explore the entirety of hilbert space in any reasonable length of time as it takes only a *polynomial number*, in n, of parameters to completely specify a local hamiltonian, but to specify a state in hilbert space requires an exponential number of parameters. Therefore, in order for a local system to explore hilbert space via its local dynamics, it ought to require at least an exponential time to visit every location. (Here we are thinking of hilbert space as a manifold rather than a linear space, i.e., we are forgetting the linear structure.) This simple counting argument can be improved [83] and it turns out that the strongest result is much better: it takes a *doubly exponential* time for the dynamics of a local hamiltonian to visit an arbitrary state in hilbert space.

To specify a pure state in the classical situation one requires only a linear (in n) number of parameters, namely, the positions of each of the particles. Another way of saying this is that for any pure classical state one can assign a definite pure state to each of the constituent particles. The main difficulty one faces in simulating a many particle quantum system is that it is difficult to describe a generic *pure* state in the hilbert space of n particles, because a general such state can be a *superposition* of possibly exponentially many classical configurations. I.e., for a general pure quantum state one cannot assign a definite pure state to each of the constituent particles as the parts do not completely determine the whole. The mismatch here is due to *quantum entanglement*, coming from the quantum correlations between the constituent particles. Thus, in order to completely describe the reduced state of a single particle we also need to keep track of the quantum correlations that develop between it and the rest of the system. This is, generically, difficult.

However, we can exploit locality structure to place bounds on how fast information can propagate through the system under natural dynamics. The physical picture to have in mind here is that of a lattice of particles connected by springs: if you wiggle one particle it should take a time proportional to the distance multiplied by a "speed of sound" for your disturbance to propagate through the system. This picture does remain metaphorically correct in the quantum setting and a fundamental result known as the *Lieb-Robinson bound* quantifies this intuition: suppose that H is the hamiltonian for a quantum spin system on a regular lattice and A and B are two observables acting nontrivially on two disjoint regions, then [61] (see [17] for a textbook treatment)

$$\|[e^{itH}Ae^{-itH}, B]\| \le c\|A\|\|B\|e^{v|t|-\kappa L}, \tag{5.1}$$

where c, v, and κ are constants, and L is the distance[4] between the regions that A and B act on. The Lieb-Robinson bound says that regardless of the initial state of the system, if Alice implements some operation at time $T = 0$ then no matter what observable Bob subsequently measures at some later time $T = t$ the amount of information that Alice can send to Bob is exponentially suppressed as the separation between them is increased [19]. Note, however, that as t increases the bound becomes exponentially weaker. This trade-off between exponentials gives rise to an effective "information propagation cone": information about Alice's operation is exponentially suppressed outside of the region $\{x \,|\, d(x, A) \le \frac{v}{\kappa}|t|\}$, where $d(x, A)$ measures the distance between site x and Alice's region.

The proof of Lieb-Robinson bound has been simplified and generalised to a variety of settings beyond the original translation invariant setting (see, e.g., [73] for a recent review): beginning with [65, 86] we've witnessed a sequence of results taking [41, 47, 71, 68, 33] it beyond the translation-invariant setting to systems with disorder [27], classical systems [87, 9], unbounded operators [28, 30, 69, 70, 84, 85], and dissipative systems

[4] Measured using the natural lattice distance.

[82, 74]. Applications of the Lieb-Robinson idea have had influences further afield, e.g., to give new proof techniques for understanding the emergence of effective equations in mean-field settings [34].

Armed with the Lieb-Robinson bound we now have the necessary quantitative tool to understand exactly how to parametrise naturally occurring states. Intuitively, although each particle does become correlated with the others after a small time, it should only be able to develop correlations with its neighbours within a distance proportional to the Lieb-Robinson velocity. One can then keep track of only these correlations and obtain a compact description of the global state. Thus we obtain an *effectively local* description of the system's state, i.e., we specify the state by recording the correlations between any given particle and only its immediate neighbours. The remarkable consequence [78, 44, 45] is that, in one dimension, there is an efficient representation known as a *matrix product state* (MPS) for the system's state, as long as $|t| \lesssim \log(N)$. An algorithm exploiting this phenomenon was already proposed in [97]; the preceding results can be interpreted as saying that this algorithm (and relatives) is indeed correct, converges, and has polynomial complexity. That is, the real-time simulation problem is in P. Further advances along these lines have shown that the simulation problem for *adiabatic evolution* of *gapped* quantum spin systems also lies within P [79].

The Lieb-Robinson technology has also been put to great effect by Hastings and coworkers in the study of the equilibrium properties of quantum lattice systems with a *spectral gap* between the lowest two eigenvalues. Beginning with the breakthrough result [41], where the the *Lieb-Schultz-Mattis* conjecture was proven in higher dimensions, much insight has been gained into the general properties of gapped quantum lattice systems (see also [72]). Firstly, the presence of the gap allows one to play the Lieb-Robinson bound off against the natural decay induced by the gap to show that all spatial ground-state correlators decay exponentially with separation [41, 42, 71, 48, 39, 13]. This result already hints at a more powerful generalisation: perhaps the entanglement entropy, as measured by the von Neumann entropy $S(\rho_L) = -\operatorname{tr}(\rho_L \log(\rho_L))$, where ρ_L is the reduced density operator of a region of length L in the ground state of a gapped spin system, is bounded by a constant? Indeed this is the case for one-dimensional systems and such an *entropy/area* law was proved in [43]. This result can be interpreted as placing the ground-state simulation problem for gapped 1D spin systems *in* NP, i.e., there exists a polynomial sized proof, written as an MPS, that the ground-state energy is low. This problem is not known to be in P.

The proof technique developed in [43] also established the key result that the ground-state of a gapped one-dimensional quantum lattice system may be efficiently approximated by an MPS. The proof is not constructive, however, and only supplies a promise that such a representation exists. It is tempting to suppose that a simple application of the variational method to MPS should find such a representation easily. However, a naive application of the variational principle to MPS runs into problems because the objective function is a nonlinear function with many local minima, and it is far from clear how to carry out the minimisation efficiently.

It turns out, however, that a practical solution to this problem has been known for many years under the guise of the *density matrix renormalisation group* (DMRG) [99]. The DMRG is now understood to be a *relaxation* of the variational method applied to MPS [88]. While the method is a heuristic, and it does suffer from local minima problems, extensive practical numerical experience has shown that these local minima do not cause any significant problems. Indeed, the DMRG is, without peer, the premiere method for the exploration of strongly correlated quantum physics in one dimension and is hard to overstate the impact of the DMRG in the field of condensed matter physics.

The question of whether there does exist a polynomial algorithm to vary over MPS has, however, been addressed. It was recently established [2, 90] that a transfer-operator type argument yields an efficient method to vary over MPS with a fixed number of variational parameters. Unfortunately, the performance of this algorithm is not comparable with the DMRG, and is unlikely to supplant the DMRG as the method of choice in the study of one-dimensional quantum systems.

The MPS variational class was proposed and studied well before the advent of hamiltonian complexity. But with the advent of quantum information theory and hamiltonian complexity new ideas began to flow into condensed matter physics leading to the development of several new variational classes for quantum lattice systems. One important example are the *projected entangled-pair states* (PEPS) [93], which are a natural generalisation of MPS to the two-dimensional setting. Another example is the *multiscale entanglement renormalisation ansatz* (MERA) [95, 96], which is particularly well-suited to the description of low-dimensional

quantum lattice systems at or near cricality. Both of these classes have been put to great effect, in combination with the variational principle, in the practical study of quantum lattice systems.

Parallel to these theoretical developments there has been the exploration of *classical* heuristics and algorithms to approximate ground-state properties of quantum spin systems. Firstly, it was realised [20] that one can exploit the low local dimensionality of quantum spins to efficiently decide the quantum 2-satisfiability problem in the spin-1/2 case. A variation on this theme exploits possible *low-degree* structure an interaction hypergraph to supply a criteria to decide whether a quantum satisfiability instance *always* has a solution. This result [8] is known as the *quantum Lovasz local lemma*. Further afield, the study of approximation algorithms and heuristics for QMA hard instances of the local hamiltonian problem has been incepted in [14, 35].

Recently, in another direction, Lieb-Robinson ideas and a technique known as *quasi-adiabatic continuation* [47, 13] have been exploited to prove several powerful results concerning the stability of topological order [22, 18]. Quasi-adiabatic continuation has been further explored in the context of disordered systems [46], where the improved bound on information flow coming from the Anderson localisation phenomena may be exploited to improve many results concerning the decay of correlations in strongly interacting quantum systems. One of the most recent triumphs of these research directions has been the proof [49] of an outstanding conjecture in mathematical physics, namely, the quantisation of the Hall conductance without averaging assumptions.

5.1. **Guiding problem: entanglement entropy/area laws.** The entanglement entropy/area law states that for a quantum system in the ground state the entropy of any subregion A should scale as the area of the boundary. This result may be regarded as an expression of the *holographic principle*, namely, that the physics in the bulk of the system can be completely described by what happens at the boundary. The existing proofs of these results have been influential in hamiltonian complexity and condensed matter physics, typically leading to many new insights and results, including new efficient methods to simulate strongly correlated systems.

Quantitative results concerning the entropy/area law are, however, difficult to obtain. Indeed, apart from special cases [66], the only settings where the result has been proved at the level of mathematical rigour pertain to gapped one-dimensional quantum spin systems and gaussian systems (see, e.g., [32] for a review). One generalisation that has proved particularly elusive is to gapped two-dimensional quantum spin systems. Here the existing proof ideas are extremely suggestive, yet just fail to yield the expected law. The difficulty is probably tied, in part, to the fact that two-dimensional quantum systems can exhibit *topological order*, which is difficult to capture using the existing proof techniques. It is clear that whatever method is developed to tackle this case will have important consequences not only for the simulation of two-dimensional quantum systems, but also for the study of exotic phases of matter.

6. Conclusion

Hamiltonian complexity, as a field, has developed rapidly in the past couple of years amassing a broad collection of foundational results exploiting a wide portfolio of techniques. Highlights include the proof that it is QMA-complete to approximate the ground-state energy of even a one-dimensional quantum lattice system and, in the other direction, proofs that the imaginary time simulation problem is in NP for *gapped* one-dimensional quantum lattice systems, and the real-time simulation problem is in P for $|t| \sim \log(n)$.

The task of finding quantum lattice systems which are hard to simulate appears to have reached a natural pause, of sorts. This is probably due to the difficulty of the proofs required and the fact that the existing techniques have been pretty much pushed as far as they can go. However, several simple-sounding problems remain, such as the commuting hamiltonian problem, and the search for inefficient quantum PCPs, and progress on these would open up many new directions.

The development of efficient algorithms for the simulation of low-dimensional quantum lattice systems continues unabated. In this direction I think the study of quantum approximation algorithms initiated in [35] holds particular promise.

One area that I am particularly excited about is that of the generalisation of the insights obtained in hamiltonian complexity to quantum systems with *continuous* degrees of freedom, including, quantum fields; we've recently understood [94, 80, 38] how to generalise the two main variational classes studied in the hamiltonian complexity literature, namely MPS and MERA, to the continuous setting, and it seems likely that many new results generalising those originating in the lattice setting will be possible.

Acknowledgements. I would like to sincerely thank Sergey Bravyi, Mick Bremner, Jens Eisert, Daniel Gottesman, Matt Hastings, Bruno Nachtergaele, Barbara Terhal, and Frank Verstraete for their careful reading of preliminary drafts of this manuscript, and their many helpful comments and suggestions. This work was supported, in part, by the cluster of excellence EXC 201 "Quantum Engineering and Space-Time Research", by the Deutsche Forschungsgemeinschaft (DFG). Thanks go also to Il Vicino, Sante Fe, where it all got started.

References

1. Scott Aaronson, *The Limits of Quantum Computers*, Sci. Am. (2008), http://www.scottaaronson.com/ writings/limitsqc-draft.pdf.
2. Dorit Aharonov, Itai Arad, and Sandy Irani, *Efficient algorithm for approximating one-dimensional ground states*, Phys. Rev. A **82** (2010), no. 1, 012315.
3. Dorit Aharonov, Itai Arad, Zeph Landau, and Umesh Vazirani, *The detectability lemma and quantum gap amplification*, Proceedings of the 41st annual ACM symposium on Theory of computing (New York, NY, USA), STOC '09, ACM, 2009, pp. 417–426.
4. ______, *Quantum Hamiltonian complexity and the detectability lemma*, arXiv:1011.3445, 2010.
5. Dorit Aharonov and Lior Eldar, *On the complexity of Commuting Local Hamiltonians, and tight conditions for Topological Order in such systems*, arXiv:1102.0770, 2011.
6. Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe, *The power of quantum systems on a line*, Comm. Math. Phys. **287** (2009), no. 1, 41–65. MR 2480741 (2010b:68058)
7. Dorit Aharonov and Tomer Naveh, *Quantum NP - A Survey*, quant-ph/0210077, 2002.
8. Andris Ambainis, Julia Kempe, and Or Sattath, *A Quantum Lovász Local Lemma*, Proceedings of the 42nd ACM symposium on Theory of computing (New York, NY, USA), STOC '10, ACM, 2010, pp. 151–160.
9. L. Amour, P. Lévy-Bruhl, and J. Nourrigat, *Dynamics and Lieb-Robinson estimates for lattices of interacting anharmonic oscillators*, Colloq. Math. **118** (2010), no. 2, 609–648. MR 2602171
10. Sanjeev Arora and Boaz Barak, *Computational complexity*, Cambridge University Press, Cambridge, 2009, A modern approach. MR 2500087 (2010i:68001)
11. Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy, *Proof verification and the hardness of approximation problems*, J. ACM **45** (1998), no. 3, 501–555. MR 1639346 (99d:68077b)
12. Sanjeev Arora and Shmuel Safra, *Probabilistic checking of proofs: a new characterization of NP*, J. ACM **45** (1998), no. 1, 70–122. MR 1614328 (99d:68077a)
13. Sven Bachmann, Spyridon Michalakis, Bruno Nachtergaele, and Robert Sims, *Automorphic Equivalence within Gapped Phases of Quantum Lattice Systems*, arXiv:1102.0842, 2011.
14. Nikhil Bansal, Sergey Bravyi, and Barbara M. Terhal, *Classical approximation schemes for the ground-state energy of quantum and classical Ising spin Hamiltonians on planar graphs*, Quantum Inf. Comput. **9** (2009), no. 7-8, 701–720. MR 2599192 (2011b:82009)
15. Francisco Barahona, *On the computational complexity of Ising spin glass models*, J. Phys. A **15** (1982), no. 10, 3241–3253. MR 684591 (84c:82022)
16. K. Binder and A. P. Young, *Spin glasses: Experimental facts, theoretical concepts, and open questions*, Rev. Mod. Phys. **58** (1986), no. 4, 801–976.
17. Ola Bratteli and Derek W. Robinson, *Operator algebras and quantum statistical mechanics. 2*, 2nd ed., Texts and Monographs in Physics, Springer-Verlag, Berlin, 1997. MR 98e:82004
18. S. Bravyi and M. B. Hastings, *A short proof of stability of topological order under local perturbations*, arXiv:1001.4363, 2010.
19. S. Bravyi, M. B. Hastings, and F. Verstraete, *Lieb-Robinson Bounds and the Generation of Correlations and Topological Quantum Order*, Phys. Rev. Lett. **97** (2006), no. 5, 050401.
20. Sergey Bravyi, *Efficient algorithm for a quantum analogue of* 2-*SAT*, quant-ph/0602108, 2006.
21. Sergey Bravyi, David P. DiVincenzo, Roberto Oliveira, and Barbara M. Terhal, *The complexity of stoquastic local Hamiltonian problems*, Quantum Inf. Comput. **8** (2008), no. 5, 361–385. MR 2451482 (2009j:81015)
22. Sergey Bravyi, Matthew B. Hastings, and Spyridon Michalakis, *Topological quantum order: Stability under local perturbations*, J. Math. Phys. **51** (2010), no. 9, 093512.
23. Sergey Bravyi, Cristopher Moore, and Alexander Russell, *Bounds on the quantum satisfiability threshold*, ICS (Andrew Chi-Chih Yao, ed.), Tsinghua University Press, 2010, pp. 482–489.
24. Sergey Bravyi and Barbara Terhal, *Complexity of Stoquastic Frustration-Free Hamiltonians*, SIAM J. on Computing **39** (2009), no. 4, 1462–1485.
25. Sergey Bravyi and Mikhail Vyalyi, *Commutative version of the local Hamiltonian problem and common eigenspace problem*, Quantum Inf. Comput. **5** (2005), no. 3, 187–215. MR 2146306 (2006i:81036)
26. Brielin Brown, Steven T. Flammia, and Norbert Schuch, *Computational Difficulty of Computing the Density of States*, arXiv:1010.3060, 2010.
27. Christian K. Burrell and Tobias J. Osborne, *Bounds on the Speed of Information Propagation in Disordered Quantum Spin Chains*, Phys. Rev. Lett. **99** (2007), 167201.
28. Paolo Buttà, Emanuele Caglioti, Sara Di Ruzza, and Carlo Marchioro, *On the propagation of a perturbation in an anharmonic system*, J. Stat. Phys. **127** (2007), no. 2, 313–325. MR 2314349 (2008m:82083)

29. H. Cohn, R. Kleinberg, B. Szegedy, and C. Umans, *Group-theoretic algorithms for matrix multiplication*, Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on, 2005, pp. 379 – 388.
30. Marcus Cramer, Alessio Serafini, and Jens Eisert, *Locality of dynamics in general harmonic quantum systems*, Quantum information and many body quantum systems, CRM Series, Ed. Norm., Pisa, 2008, pp. 51–73. MR 2544074 (2011c:82010)
31. Irit Dinur, *The PCP theorem by gap amplification*, J. ACM **54** (2007), no. 3, Art. 12, 44. MR 2314254 (2008f:68037b)
32. J. Eisert, M. Cramer, and M. B. Plenio, *Colloquium: Area laws for the entanglement entropy*, Rev. Mod. Phys. **82** (2010), no. 1, 277–306.
33. Jens Eisert and Tobias J. Osborne, *General Entanglement Scaling Laws from Time Evolution*, Phys. Rev. Lett. **97** (2006), no. 15, 150404.
34. László Erdős and Benjamin Schlein, *Quantum dynamics with mean field interactions: a new approach*, J. Stat. Phys. **134** (2009), no. 5-6, 859–870. MR 2518972 (2010f:81318)
35. Sevag Gharibian and Julia Kempe, *Approximation algorithms for QMA-complete problems*, arXiv:1101.3884, 2011.
36. Daniel Gottesman and Sandy Irani, *The quantum and classical complexity of translationally invariant tiling and hamiltonian problems*, Proceedings of the 2009 50th Annual IEEE Symposium on Foundations of Computer Science (Washington, DC, USA), FOCS '09, IEEE Computer Society, 2009, pp. 95–104.
37. Daniel Gottesman, *private communication*, 2011.
38. Jutho Haegeman, Tobias J. Osborne, Henri Verschelde, and Frank Verstraete, *Entanglement renormalization for quantum fields*, arXiv:1102.5524, 2011.
39. Eman Hamza, Spyridon Michalakis, Bruno Nachtergaele, and Robert Sims, *Approximating the ground state of gapped quantum spin systems*, J. Math. Phys. **50** (2009), no. 9, 095213. MR 2566893 (2011d:82012)
40. Johan Håstad, *Some optimal inapproximability results*, J. ACM **48** (2001), no. 4, 798–859 (electronic). MR 2144931 (2006c:68066)
41. M. B. Hastings, *Lieb-Schultz-Mattis in higher dimensions*, **69** (2004), no. 10, 104431.
42. ———, *Locality in Quantum and Markov Dynamics on Lattices and Networks*, (2004), no. 2, 4.
43. ———, *An area law for one-dimensional quantum systems*, J. Stat. Mech. (2007), no. 8, P08024. MR 2338267 (2008h:82010)
44. ———, *Observations outside the light cone: Algorithms for nonequilibrium and thermal states*, Phys. Rev. B **77** (2008), no. 14, 144302.
45. ———, *Light-cone matrix product*, J. Math. Phys. **50** (2009), no. 9, 095207, 17. MR 2566887 (2011b:82079)
46. M. B. Hastings, *Quasi-adiabatic Continuation for Disordered Systems: Applications to Correlations, Lieb-Schultz-Mattis, and Hall Conductance*, arXiv:1001.5280, 2010.
47. M. B. Hastings and Xiao-Gang Wen, *Quasiadiabatic continuation of quantum states: The stability of topological ground-state degeneracy and emergent gauge invariance*, Phys. Rev. B **72** (2005), no. 4, 045141.
48. Matthew B. Hastings and Tohru Koma, *Spectral Gap and Exponential Decay of Correlations*, Comm. Math. Phys. **265** (2006), no. 3, 781–804.
49. Matthew B. Hastings and Spyridon Michalakis, *Quantization of Hall Conductance For Interacting Electrons Without Averaging Assumptions*, arXiv:0911.4706, 2009.
50. Dominik Janzing, Pawel Wocjan, and Thomas Beth, *Cooling and Low Energy State Preparation for 3-local Hamiltonians are FQMA-complete*, quant-ph/0303186, 2003.
51. Mark Jerrum and Alistair Sinclair, *Polynomial-time approximation algorithms for the Ising model*, SIAM J. Comput. **22** (1993), no. 5, 1087–1116. MR 94g:82007
52. Stephen P. Jordan, David Gosset, and Peter J. Love, *Quantum-Merlin-Arthur–complete problems for stoquastic Hamiltonians and Markov matrices*, Phys. Rev. A **81** (2010), no. 3, 032331.
53. Alastair Kay, *Quantum-Merlin-Arthur-complete translationally invariant hamiltonian problem and the complexity of finding ground-state energies in physical systems*, Phys. Rev. A **76** (2007), no. 3, 030307(R).
54. Julia Kempe, Alexei Kitaev, and Oded Regev, *The complexity of the local Hamiltonian problem*, FSTTCS 2004: Foundations of software technology and theoretical computer science, Lecture Notes in Comput. Sci., vol. 3328, Springer, Berlin, 2004, pp. 372–383. MR 2140890
55. Julia Kempe and Oded Regev, *3-local Hamiltonian is QMA-complete*, Quantum Inf. Comput. **3** (2003), no. 3, 258–264. MR 1985547 (2005b:68106)
56. Scott Kirkpatrick and Bart Selman, *Critical Behavior in the Satisfiability of Random Boolean Expressions*, Science **264** (1994), no. 5163, 1297–1301.
57. A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and quantum computation*, Graduate Studies in Mathematics, vol. 47, American Mathematical Society, Providence, RI, 2002. MR 1 907 291
58. C. R. Laumann, A. M. Läuchli, R. Moessner, A. Scardicchio, and S. L. Sondhi, *Product, generic, and random generic quantum satisfiability*, Phys. Rev. A **81** (2010), no. 6, 062345.
59. C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi, *Random quantum satisfiability*, Quantum Inf. Comput. **10** (2010), no. 1-2, 1–15.
60. ———, *Statistical mechanics of classical and quantum computational complexity*, arXiv:1009.1635, 2010.
61. Elliott H. Lieb and Derek W. Robinson, *The finite group velocity of quantum spin systems*, Comm. math. Phys. **28** (1972), 251–257. MR 47 #1415
62. Yi-Kai Liu, *Consistency of local density matrices is QMA-complete*, Approximation, randomization and combinatorial optimization, Lecture Notes in Comput. Sci., vol. 4110, Springer, Berlin, 2006, pp. 438–449. MR 2305030 (2008a:68081)
63. ———, *The Local Consistency Problem for Stoquastic and 1-D Quantum Systems*, arXiv:0712.1388, 2007.

64. Yi-Kai Liu, Matthias Christandl, and F. Verstraete, *Quantum Computational Complexity of the N-Representability Problem: QMA Complete*, Phys. Rev. Lett. **98** (2007), no. 11, 110503.
65. C. Marchioro, A. Pellegrinotti, M. Pulvirenti, and L. Triolo, *Velocity of a perturbation in infinite lattice systems*, J. Statist. Phys. **19** (1978), no. 5, 499–510. MR 515469 (80b:82053)
66. Lluís Masanes, *Area law for the entropy of low-energy states*, Phys. Rev. A **80** (2009), no. 5, 052104.
67. Ramis Movassagh, Edward Farhi, Jeffrey Goldstone, Daniel Nagaj, Tobias J. Osborne, and Peter W. Shor, *Unfrustrated qudit chains and their ground states*, Phys. Rev. A **82** (2010), 012318.
68. Bruno Nachtergaele, Yoshiko Ogata, and Robert Sims, *Propagation of correlations in quantum lattice systems*, J. Stat. Phys. **124** (2006), no. 1, 1–13. MR 2256615 (2007k:82075)
69. Bruno Nachtergaele, Hillel Raz, Benjamin Schlein, and Robert Sims, *Lieb-Robinson bounds for harmonic and anharmonic lattice systems*, Comm. Math. Phys. **286** (2009), no. 3, 1073–1098. MR 2472026 (2010k:82016)
70. Bruno Nachtergaele, Benjamin Schlein, Robert Sims, Shannon Starr, and Valentin Zagrebnov, *On the existence of the dynamics for anharmonic quantum oscillator systems*, Rev. Math. Phys. **22** (2010), no. 2, 207–231. MR 2655205
71. Bruno Nachtergaele and Robert Sims, *Lieb-Robinson Bounds and the Exponential Clustering Theorem*, Comm. Math. Phys. **265** (2006), no. 1, 119–130.
72. ———, *A multi-dimensional Lieb-Schultz-Mattis theorem*, Comm. Math. Phys. **276** (2007), no. 2, 437–472. MR 2346396 (2008i:82015)
73. ———, *Lieb-Robinson bounds in quantum many-body physics*, Entropy and the quantum, Contemp. Math., vol. 529, Amer. Math. Soc., Providence, RI, 2010, pp. 141–176. MR 2681770
74. Bruno Nachtergaele, Anna Vershynina, and Valentin A. Zagrebnov, *Lieb-Robinson Bounds and Existence of the Thermodynamic Limit for a Class of Irreversible Quantum Dynamics*, arXiv:1103.1122, 2011.
75. Daniel Nagaj and Shay Mozes, *New construction for a QMA complete three-local Hamiltonian*, J. Math. Phys. **48** (2007), no. 7, 072104, 13. MR 2337664 (2008g:81047)
76. Michael A. Nielsen and Isaac L. Chuang, *Quantum computation and quantum information*, Cambridge University Press, Cambridge, 2000. MR 1 796 805
77. Roberto Oliveira and Barbara M. Terhal, *The complexity of quantum spin systems on a two-dimensional square lattice*, Quantum Inf. Comput. **8** (2008), no. 10, 900–924. MR 2522687 (2011d:81077)
78. Tobias J. Osborne, *Efficient Approximation of the Dynamics of One-Dimensional Quantum Spin Systems*, Phys. Rev. Lett. **97** (2006), no. 15, 157202.
79. ———, *Simulating adiabatic evolution of gapped spin systems*, Phys. Rev. A **75** (2007), no. 3, 032321.
80. Tobias J. Osborne, Jens Eisert, and Frank Verstraete, *Holographic quantum states*, Phys. Rev. Lett. **105** (2010), no. 26, 260401.
81. Christos H. Papadimitriou, *Computational complexity*, Addison-Wesley Publishing Company, Reading, MA, 1994. MR 1251285 (95f:68082)
82. David Poulin, *Lieb-Robinson bound and locality for general Markovian quantum dynamics*, Phys. Rev. Lett. **104** (2010), no. 19, 190401, 4. MR 2647323
83. David Poulin, Angie Qarry, R. D. Somma, and Frank Verstraete, *Quantum simulation of time-dependent Hamiltonians and the convenient illusion of Hilbert space*, arXiv:1102.1360, 2011.
84. Isabeau Prémont-Schwarz, Alioscia Hamma, Israel Klich, and Fotini Markopoulou-Kalamara, *Lieb-robinson bounds for commutator-bounded operators*, Phys. Rev. A **81** (2010), no. 4, 040102.
85. Isabeau Prémont-Schwarz and Jeff Hnybida, *Lieb-robinson bounds on the speed of information propagation*, Phys. Rev. A **81** (2010), no. 6, 062107.
86. Charles Radin, *Signal propagation in lattice models of quantum many-body systems*, Comm. Math. Phys. **62** (1978), no. 2, 159–166. MR 0503321 (58 #20103)
87. Hillel Raz and Robert Sims, *Estimating the Lieb-Robinson velocity for classical anharmonic lattice systems*, J. Stat. Phys. **137** (2009), no. 1, 79–108. MR 2556737
88. Ulrich Schollwöck, *The density-matrix renormalization group in the age of matrix product states*, Ann. Phys. **326** (2011), no. 1, 96 – 192.
89. Norbert Schuch, *Complexity of commuting hamiltonians on a square lattice of qubits*, arXiv:1105.2843, 2011.
90. Norbert Schuch and J. Ignacio Cirac, *Matrix product state and mean-field solutions for one-dimensional systems can be found efficiently*, Phys. Rev. A **82** (2010), no. 1, 012314.
91. Norbert Schuch and Frank Verstraete, *Computational complexity of interacting electrons and fundamental limitations of density functional theory*, Nat. Phys. **5** (2009), no. 10, 732 – 735.
92. Yaoyun Shi and Shengyu Zhang, *Note on quantum counting classes*, http://www.cse.cuhk.edu.hk/%7Esyzhang/papers/SharpBQP.pdf, 2009.
93. F. Verstraete and J. I. Cirac, *Renormalization algorithms for Quantum-Many Body Systems in two and higher dimensions*, 2004.
94. ———, *Continuous matrix product states for quantum fields*, Phys. Rev. Lett. **104** (2010), no. 19, 190405.
95. G. Vidal, *A class of quantum many-body states that can be efficiently simulated*, cond-mat/0610099, 2006.
96. ———, *Entanglement Renormalization*, Phys. Rev. Lett. **99** (2007), no. 22, 220405.
97. Guifré Vidal, *Efficient Simulation of One-Dimensional Quantum Many-Body Systems*, Phys. Rev. Lett. **93** (2003), no. 4, 040502.
98. Tzu-Chieh Wei, Michele Mosca, and Ashwin Nayak, *Interacting Boson Problems Can Be QMA Hard*, Phys. Rev. Lett. **104** (2010), no. 4, 040501.

99. S. R. White, *Density matrix formulation for quantum renormalization groups*, Phys. Rev. Lett. **69** (1992), no. 10, 2863–2866.

E-mail address: tobias.osborne@rhul.ac.uk